

Supporting Information

Physical Cryptographic Verification of Nuclear Warheads

R. Scott Kemp Areg Danagoulian Ruairidh R. Macdonald Jayson R. Vavrek

Contents

1	Nuclear Resonance Fluorescence (NRF)	2
1.1	Interaction cross section	2
1.2	NRF as an injective function assuring soundness	2
2	Single Pixel X-Ray Transform (\mathcal{K})	3
2.1	Relationship to the physical measurement	4
2.2	Proof of the global uniqueness of the \mathcal{K} transform on \mathbb{R}^n	5
3	System Overview	5
3.1	Concept of operations	5
3.2	Physical components of a transmission-NRF measurement	6
4	Derivation of the Expected NRF Signal	7
4.1	Uniform object and uniform foil	7
4.2	Heterogenous device	9
4.3	Non-invertibility of single-energy measurements	9
5	Monte Carlo Simulations	9
5.1	Bremsstrahlung interrogation beam	9
5.2	Simulation geometries	11
5.3	Monte Carlo results	12
5.4	Extrapolation to fixed measurement times	17
6	Statistical Tests of Identicality	18
6.1	Statistics of single projection comparisons	18
6.2	Statistics of sampling multiple projections	18
6.2.1	Definition of a geometric hoax	19
6.2.2	A discrete approximation for estimating the detectability of a geometric hoax	19
7	Information Security	21
7.1	Information security in real systems	22
7.1.1	X-ray fluorescence	22
7.1.2	Nonmonochromatic NRF measurements	22
7.1.3	Prior assumption of geometries in estimating mass	23
8	Operational Considerations	24
8.1	The measured object	24
8.2	Dose	24
8.3	Electronics and count rates	25
8.4	System integrity	26

1 Nuclear Resonance Fluorescence (NRF)

The key physical process at the core of the proposed verification system is nuclear resonance fluorescence (NRF). NRF entails the collective excitation of a nucleus caused by the absorption of an external real photon. After a short time ($T_{1/2} \ll \text{ns}$) the nucleus decays to a lower energy state, emitting a photon with energy equal to the difference of the initial and final nuclear states. In this regard NRF is the nuclear analogue of atomic fluorescence, but at $\sim \text{MeV}$ energy scales.

A nucleus excited by an NRF interaction can be described by two primary modes of oscillation. The first mode involves neutron and proton groups vibrating around the common center in scissor-like oscillations [1]. The scissor mode amounts to a magnetic dipole excitation. The second mode of excitation involves a combination of octupole and quadrupole oscillations, which amount to an electric dipole excitation [2].

1.1 Interaction cross section

The cross section $\sigma_{\text{NRF}}(E)$ describing the probability that an incoming photon excites an NRF resonance, where the energy E is in the center-of-mass frame, is described by a Breit-Wigner distribution:

$$\sigma_{\text{NRF}}(E) = \frac{2J+1}{2J_0+1} \frac{\pi}{2} \frac{(\hbar c)^2}{E_r^2} \frac{\Gamma \Gamma_0}{(E - E_r)^2 + (\Gamma/2)^2} \quad (\text{S.1})$$

where J and J_0 are the angular momentum quantum numbers of the excited and ground states, respectively, E_r is the energy of the excited state, Γ_0 is the width of the ground state, and Γ is a sum over the widths of all possible transition states [3]. For the most common NRF states the widths are $\mathcal{O}(\text{meV})$.

The thermal motion of atoms Doppler broadens the resonance, which can be approximated as a normal distributions with a width of $\Delta = E_c \sqrt{k_B T / M} \sim \mathcal{O}(\text{eV})$. Here T is the temperature¹, and M is the mass of the nucleus in a system of units where $c \equiv 1$ [3]. Due to the narrow width of the NRF cross sections, direct photon scattering experiments can only reveal the total energy-integrated cross sections and the resonant energy of the state. For a resonance where the excitation always decays to the ground state ($\Gamma = \Gamma_0$) the spin states and energy widths are related to the observables by [3]

$$\int \sigma_{\text{NRF}} dE = \frac{2J_1+1}{2J_0+1} \pi^2 \hbar^2 \Gamma / E_r^2 \quad (\text{S.2})$$

The above energy-integrated cross section is conserved both after Doppler widening as well as after the effects of detection instrumentation's finite energy resolution. For a high purity germanium (HPGe) detector, the energy resolution is $\mathcal{O}(\text{keV})$. A number of experiments have taken place to measure the integrated cross sections for ^{238}U [4], ^{235}U and ^{239}Pu [5], and ^{240}Pu [6]. Figure S.1 shows the observed energy spectrum from a highly enriched uranium (HEU) sample, showing the NRF lines corresponding to ^{235}U . The cited measurements are used in all analytical estimates and Monte Carlo simulations of the NRF-based verification system.

1.2 NRF as an injective function assuring soundness

The main feature of NRF for our purposes is the sparsity and isotopic uniqueness of the resonant energies of the nucleus. The isotopic uniqueness arises because the quantized energy levels allowed for the nucleus are specific to the composition of the nucleus and thus its identity. The sparsity arises because of the very narrow energy width of the resonances, and are enhanced by quantum-mechanical selection rules. The sparsity holds even when the apparent width of lines are broadened by $\mathcal{O}(10^3)$ times because of the limited energy resolution of a practical detector. For example, the average ratio of line widths to line spacing for an HPGe-measured spectrum is $\mathcal{O}(0.01)$. These two effects combined render NRF an exceptionally useful tool for

¹For greater accuracy in solids, T can be replaced by the effective temperature

$$T_{\text{eff}} = 3T \left(\frac{T}{\theta} \right)^3 \int_0^{\theta/T} x^3 \left(\frac{1}{e^x - 1} + \frac{1}{2} \right) dx,$$

where θ is the Debye temperature of the target material. This accounts for the fact that the energy distribution of nuclei in a solid does not follow a perfect Maxwell-Boltzmann distribution due to the solid's lattice structure.




Figure S.1: NRF lines from a sample of sample of highly enriched uranium. The resonances for ^{235}U are indicated with green arrows [5]. The energy spectrum from the target holder was measured separately (black), allowing the subtraction of its contributions from the total spectrum (red). The line widths are almost entirely determined by the resolution of the high purity germanium (HPGe) detectors and $\mathcal{O}(\text{keV})$.

ensuring that a substitution of one material for another will be detected. In mathematical terms, NRF acts as an injective map, which is a necessary condition for the proof system to be *sound*. By comparison, prior verification work using physical systems to protect against the disclosure of information have used neutron scattering [7]. However, neutron scattering cross sections are non-sparse. They can often be extremely similar between isotopes, for example between weapons-grade plutonium (WGPu) and inexpensive depleted uranium. As a result, neutron scattering is relatively insensitive to material substitutions.

2 Single Pixel X-Ray Transform (\mathcal{K})

The key theoretical construct underpinning the soundness of the system is the ability to make tomographic measures of the candidate object. All tomographic methods rely on invertible integral-geometric transforms. Transmission radiography, which is similar to the NRF apparatus described in Section 3.2, uses the X-Ray Transform (\mathcal{P}). In three dimensions, the X-Ray Transform \mathcal{P} of f at (r, θ) is the line integral of f over the straight lines through the point $r \in \mathbb{R}^3$ with direction $\theta \in \mathbb{S}^2$. In our case, however, we cannot allow for the detector to have reconstructable angular resolution in the projection space \mathbb{S}^2 , either through pixelated detectors as is typically done in transmission x-ray tomography, or through single-pixel coded apertures, because relationships in the \mathbb{S}^2 projection space contain potentially sensitive information about the warhead’s geometry.

The Single Pixel X-ray Transform (\mathcal{K}) proposed here is an integral-geometric transform $\mathcal{K} : \mathbb{R}^n \rightarrow \mathbb{R}$. A straightforward collapse of dimensions through multiple integration would normally not be uniquely invertible because of the Fubini theorem. However, by using a nonlinear monotonic weighting, \mathcal{K} is able to provide a mapping with global uniqueness.

2.1 Relationship to the physical measurement

Let f be a continuous real-valued function of compact support on the n -dimensional Euclidian space \mathbb{R}^n . Let $l = r + s\theta$ be a line passing through the x-ray source point r in the direction of the unit vector $\theta \in \mathbb{S}^{n-1}$, where \mathbb{S}^{n-1} is the projective space. The conventional *X-Ray Transform* is

$$\mathcal{P}f(l) = \int f(r + s\theta) ds \quad (\text{S.3})$$

where the notation \mathcal{P} is treated as if it were an operator. The transform is defined for each individual line l and the integral is Lebesgue. In this section we want to define the *Single Pixel X-ray Transform* which Lebesgue-integrates over the set of all lines $L(r) : \{l(r, s\theta)\}$ passing through a point r , as illustrated in Figure S.2.




Figure S.2: Illustration of the Single Pixel X-Ray Transform's operation. Instead of a pixelated detector arrangement required for the conventional X-Ray transform, the Single-Pixel X-Ray Transform uses a single detector that integrates all rays in the projection space.

Let the Single Pixel X-ray Transform (\mathcal{K}) be defined as

$$\mathcal{K}f(r) \equiv \int_{L(r)} e^{-\mathcal{P}f(l)} d\theta = \int_{\mathbb{S}^{n-1}} e^{-\int f(r+s\theta) ds} d\theta \quad (\text{S.4})$$

where again the symbol \mathcal{K} is treated as a operator and the Lebesgue integral is over all directions.

In practical terms, the nonlinear relation between $\mathcal{K}f$ and $\mathcal{P}f$ assures that, if f is a density map, the value $\mathcal{K}f$ is not simply the invariant mass. The exponential function is chosen for the nonlinear function because (1) it is monotonic and (2) NRF attenuation is naturally exponential in space. This holds not only for the parallel-beam approximation used in Equation S.5, but also for the more general fan-beam case in \mathbb{R}^3 .

Proof. Let $\{r = \|\vec{r}\| \in [0, +\infty), \vartheta \in [0, 2\pi), \varphi \in [0, \pi]\}$ be a spherical coordinate system centered on the x-ray source. Note: φ is a coordinate and ϕ is flux (scalar); \mathbf{J}_{beam} is the photon current (vector) at some point (r, ϑ, φ) . From the continuity equation, we have:

$$\begin{aligned} \nabla \cdot \mathbf{J}_{\text{beam}} &= -\phi f \\ \frac{1}{r^2} \frac{d[r^2(\mathbf{J} \cdot \hat{\mathbf{r}})]}{dr} &= -\phi_r f \\ \frac{1}{r^2} \frac{d(r^2 \phi_r)}{dr} &= -\phi_r f \\ \frac{1}{r^2 \phi_r} \frac{d(r^2 \phi_r)}{dr} &= -f \\ \ln(r^2 \phi_r) + C &= -\int_r f dr \end{aligned}$$

The initial condition is the flux from the unattenuated beam: $\phi_o(\vec{r}, E, \hat{\Omega}) = \phi_o(\vec{r}, E, \hat{\mathbf{r}}) = [\Phi_o / (4\pi r^2)] dE d^2\Omega$,

which gives $C = -\ln(r^2\phi_o)$, and:

$$\begin{aligned}\ln\left(\frac{\phi_r}{\phi_o}\right) &= -\int_r f dr \\ \left(\frac{\phi_r}{\phi_o}\right) &= e^{-\int_r f dr} = e^{-\mathcal{P}f} \quad (\text{by Eq. S.3}) \\ \iint_{\text{detector}} \left(\frac{\phi_r}{\phi_o}\right) d\vartheta d\varphi &= \int_{\mathbb{S}^2} e^{-\mathcal{P}f} d\theta = \mathcal{K}f \quad \blacksquare\end{aligned}$$

2.2 Proof of the global uniqueness of the \mathcal{K} transform on \mathbb{R}^n

The \mathcal{K} transform exhibits global uniqueness, which is to say that given the set of all projections for every r , the transform will provide a unique, injective mapping of the geometry of the warhead. The authors thank Gunther Uhlmann for developing this proof.

Let $\mathcal{P}f$ be the line integral (X-Ray) transform on \mathbb{R}^n defined by Equation S.3 and let \mathcal{K} be the non-linear operator defined by Equation S.4. Let L^1 be the space of all Lebesgue integrable functions on \mathbb{R}^n , and L^∞ be the essentially bounded measurable functions of the dual space. The following result says that $-\mathcal{K}$ is a monotone operator.

Theorem 1. *For $f, g \in L^1$ such that $\mathcal{K}f, \mathcal{K}g \in L^\infty$ then*

$$\langle \mathcal{K}f - \mathcal{K}g, f - g \rangle \leq 0,$$

and the inequality is strict when $f \neq g$.

Proof. Note that $\mathcal{P}f(r, \theta)$ is constant on lines in direction θ , so $\mathcal{P}f(r, \theta) = \mathcal{P}f(E_\theta r, \theta)$, where E_θ denotes orthogonal projection on θ^\perp . Therefore:

$$\begin{aligned}\langle \mathcal{K}f - \mathcal{K}g, f - g \rangle &= \int_{\mathbb{R}^n} \int_{\mathbb{S}^{n-1}} \left(e^{-\mathcal{P}f(r, \theta)} - e^{-\mathcal{P}g(r, \theta)} \right) (f - g)(r) d\theta dr \\ &= \int_{\mathbb{S}^{n-1}} \int_{\mathbb{R}^n} \left(e^{-\mathcal{P}f(r, \theta)} - e^{-\mathcal{P}g(r, \theta)} \right) (f - g)(r) dr d\theta \\ &= \int_{\mathbb{S}^{n-1}} \int_{\theta^\perp} \left(e^{-\mathcal{P}f(E_\theta r, \theta)} - e^{-\mathcal{P}g(E_\theta r, \theta)} \right) \int_{\mathbb{R}} (f - g)(E_\theta r + s\theta) ds dr_H d\theta \\ &= \int_{\mathbb{S}^{n-1}} \int_{\theta^\perp} \left(e^{-\mathcal{P}f(E_\theta r, \theta)} - e^{-\mathcal{P}g(E_\theta r, \theta)} \right) (\mathcal{P}f(E_\theta r, \theta) - \mathcal{P}g(E_\theta r, \theta)) dr_H d\theta\end{aligned}$$

where dr_H is Lebesgue measure on the hyperplane θ^\perp . The integrand has the form $(e^{-s} - e^{-t})(s - t)$, which is negative except when $s = t$ and so $\langle \mathcal{K}f - \mathcal{K}g, f - g \rangle < 0$ unless $\mathcal{P}f = \mathcal{P}g$ almost everywhere. Then uniqueness for the X-Ray transform [8] implies that $g = f$ almost everywhere. \blacksquare

3 System Overview

3.1 Concept of operations

In order to verify a nuclear warhead within the requirements of *soundness*, *completeness*, and *information protection*, the measurement needs to be embedded in a sequence of steps. In this framework there are two parties: the representatives of the country which is to prove compliance with the arms-reduction treaty (from now referred to as the *host*); and the inspection team from the treaty partner country (from now referred to as the *inspectors*). The verification protocol consists of the following steps:

1. The inspectors arrive to the host country. A template warhead is established by a mechanism as described in the main text. Several such template warheads can be chosen and compared to each other

to increase the confidence that the warheads selected are a true representative sample of all warheads deployed. National technical means (e.g., satellite surveillance) can be used to confirm that warheads on delivery systems remained undisturbed prior to and after the template selection process, further enhancing template confidence.

2. Chain of custody mechanisms, such as unique-ID tags, seals, and surveillance, are used to ensure the selected template is not modified while being transported to the inspection facility.
3. The host present a set of candidate warheads for dismantlement. Chain of custody procedures might also be used to help ensure these warheads came from delivery vehicles, although this is optional. The candidate warheads enter joint custody of the inspectors and host at the inspection facility.
4. One by one, each candidate weapon is compared to the template in the provably sound, zero-disclosure test outlined in this article. Comparisons are made projection by projection. A single projection follows these steps:
 - (a) Projection parameters are randomly selected, where a projection consists of an angle of rotation for the warhead and a translation distance (Section 2).
 - (b) An certified encrypting foil is selected and installed into the specialized x-ray apparatus (Section 3.2).
 - (c) A transmission-NRF measurement of the template is performed at the selected projection and with the selected foil.
 - (d) A transmission-NRF measurement of the candidate is performed at the same projection parameters and with the same foil.
 - (e) If the measured signal (Section 4) of the candidate matches within allowed projection statistics (Section 6.1) the measured signal of the template, it passes the projection.
5. Depending on the overall statistical confidence required (Section 6.2), additional randomly selected projections may be taken by repeating step 4. In the simulations presented here, a set of canonical hoaxes were rejected with $> 98\%$ probability using only one projection.

Once these steps are completed, the now-authenticated weapon is sent for dismantlement, and its disposition is counted towards the host country's obligations in the arms-reduction treaty.

The secrecy of the warhead design is protected by the one-time-pad encryption effect of the foil, which is manufactured by the host country with a composition known only to the host. The host can vary the areal density of the foil by several orders of magnitude, in practice, from microns to centimeters. The integrated density of the warhead after a projection measurement will therefore be uncertain to the same extent. A uniform foil would be self validating. Under the assumption of uniformity, the inspector needs to only ascertain that it contains the agreed upon isotopes (e.g. ^{235}U , ^{239}Pu , ^{238}U , ^{14}N , ^{12}C , etc.) and the very observation of NRF spectral lines specific to those isotopes is sufficient to validate that the foil contains the requisite isotope. Thereafter, data must be collected until the statistical-confidence goal is met. If the foil is too thin, or lacks entirely one isotope, it will not be possible to meet the goal in the agreed measurement time and the inspector can reject the foil as inadequate. However, uniformity in the foil would need to be established by some mechanism. Additionally, non-uniform foils may be desirable to further prevent statistical estimates of the foil's parameters (Section 7).

3.2 Physical components of a transmission-NRF measurement

The measurement of a projection involves determining the NRF attenuation by the warhead in an information-secure way. This measurement is effected by an apparatus that illuminates the warhead (and its container) with a beam of photons (either a continuous bremsstrahlung or tuned monoenergetic source) in the 2–9 MeV range. The interaction proceeds as illustrated in Figure S.3.

- The incident flux, $\phi_0(E) \equiv dN/dE$, illuminates the warhead and its container, which are collectively called an *object*.

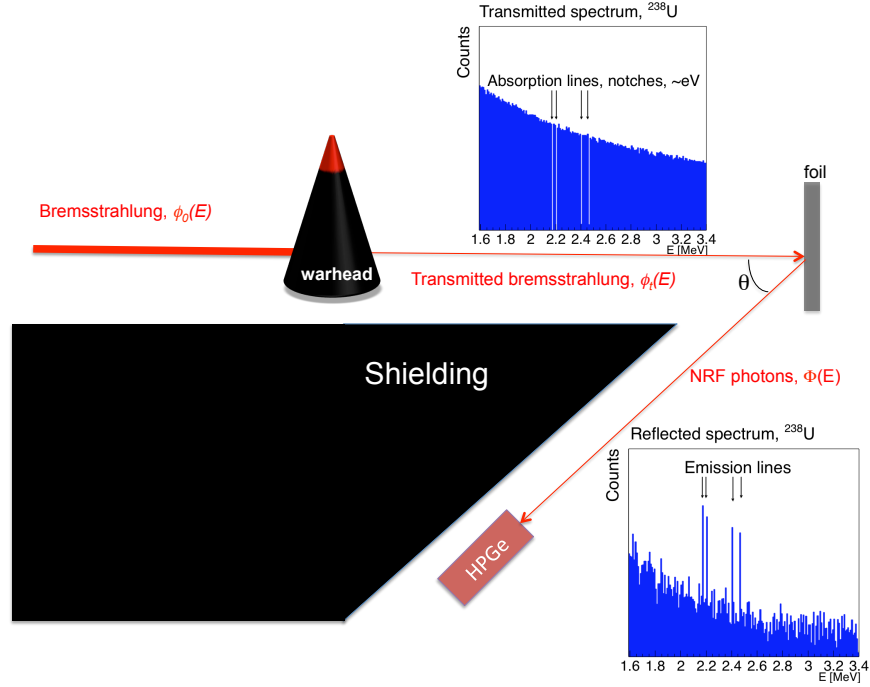


Figure S.3: Schematic drawing of the system, showing the various components.

- The object is positioned and oriented according to the projection parameters selected (Section 3.1, step 4). NRF interactions absorb photons from the beam at resonant energies (Section 1). Excited nuclei immediately decay, but the relaxation photons are emitted into a heavily shielded room so that no detector is able to measure them.
- The transmitted beam exiting the object is $\phi_t(E)$. This beam is the modified $\phi_0(E)$, attenuated both by NRF as well as non-resonant interactions. There is no known way to directly measure the NRF attenuation in this beam, but to be sure that no measurements are taken directly, the beam passes into a second shielded area to ensure that it cannot be measured directly.
- The beam falls incident on the foil. Each isotope of interest is present in the foil with some areal density X , which is the product of density and the foil thickness. The composition of the foil is unknown to the inspector.
- The detector is placed at angle θ to the beam incident on the foil to measure the NRF decay photons from the foil, which emerge with spectral exitance (photons per unit surface area per unit energy) M . The estimated signal measured in the detector is derived in Section 4.

4 Derivation of the Expected NRF Signal

4.1 Uniform object and uniform foil

Consider the simplified case of a uniform object of one isotope and a uniform foil of the same isotope. The transmitted spectrum $\phi_t(E)$ is determined from

$$\phi_t(E) = \phi_0(E) \cdot \exp \left(-D \frac{N_A}{A} (\sigma_{\text{NRF}}(E) + \sigma_{\text{nr}}(E)) \right) \quad (\text{S.5})$$

where N_A is Avogadro's number, A is the atomic weight, σ_{NRF} is the NRF absorption total cross section at the NRF energy of interest, while σ_{nr} is the sum of the cross sections of all non-resonant interactions. It can be written as a sum of Compton scattering, pair production, photoelectric effect, Rayleigh scattering, as well as the much weaker nuclear Compton scattering and Delbrück scattering:

$$\sigma_{nr} = Z\sigma_{CS} + \sigma_{pp} + \sigma_{pe} + \sigma_{RS} + \sigma_{nCS} + \sigma_D. \quad (\text{S.6})$$

Here it is convenient to perform the following substitution, using the definitions of mass attenuation coefficients: $\frac{N_A}{A}\sigma_{\text{NRF}} \equiv \mu_{\text{NRF}}$, and $\frac{N_A}{A}\sigma_{nr} \equiv \mu_{nr}$, yielding

$$\phi_t(E) = \phi_0(E) \cdot \exp(-D(\mu_{\text{NRF}}(E) + \mu_{nr}(E))). \quad (\text{S.7})$$

In the system shown in Figure S.3, $\phi_t(E)$ is the spectrum of the photon flux incident on the foil. This allows the determination of the NRF signal produced by the foil. As the incident flux traverses the foil, for an infinitesimal areal thickness dx at x , the number of NRF photons emitted from the foil at energy E depends on the cross section, resulting in the following equation:

$$dn = dx \phi_t(E) \cdot \frac{N_A}{A} \exp \left[-x \left(\mu_{\text{NRF}} + \mu_{nr} \frac{1 + \cos \theta}{\cos \theta} \right) \right] \left[\sigma_{\text{NRF}} \frac{\Omega}{4\pi} \right] \quad (\text{S.8})$$

where θ is the angle of the detector relative to the normal of the foil, Ω is the solid angle of the detector, A is the atomic mass and N_A is Avogadro's number. The exponential term accounts for non-resonant attenuation into and out of the foil. Several approximations are used in this derivation for the purposes of clarity.²

Using these approximations, the prior definition of $\mu_{\text{NRF}}(E)$, and integrating for x , the following expression for total NRF counts in the detector can be obtained:

$$\begin{aligned} \Phi_{\text{NRF}}(E) &= \int_0^X \frac{dn}{dx} dx \\ &= \phi_0(E) \exp(-D(\mu_{\text{NRF}} + \mu_{nr})) \frac{\mu_{\text{NRF}}}{\mu_{\text{NRF}} + \mu_{nr} \frac{1 + \cos \theta}{\cos \theta}} \times \\ &\quad \left[1 - \exp \left(-X \left(\mu_{\text{NRF}} + \mu_{nr} \frac{1 + \cos \theta}{\cos \theta} \right) \right) \right] \times \frac{\Omega}{4\pi} \end{aligned} \quad (\text{S.9})$$

where all the attenuation coefficients are functions of photon energy. The equation can then be rearranged into

$$\Phi_{\text{NRF}}(E) = A_{nr}(E) \exp(-\mathbf{D}\mu_{\text{NRF}}) \times [1 - \exp(\mathbf{T}\mu_{\text{NRF}})] \Omega/4\pi \quad (\text{S.10a})$$

where

$$A_{nr} \equiv \phi_0(E) \mu_{\text{NRF}} \exp(-\mu_{nr} \mathbf{D}) / (\mu_{\text{NRF}} + \mu_{nr}(1 + \cos \theta) / \cos \theta) \quad (\text{S.10b})$$

$$\mathbf{T} \equiv \mathbf{X} \left(1 + \frac{1 + \cos \theta}{\cos \theta} \mu_{nr} / \mu_{\text{NRF}} \right) \quad (\text{S.10c})$$

Here A_{nr} is proportional to the incident flux which is non-resonantly attenuated by the object, and \mathbf{T} is the *effective thickness* of the foil. The spectral exitance given in the main text is $M = \Phi_{\text{NRF}}|_{\Omega=4\pi}$.

²To simplify the mathematical formalism, here the NRF cross section is approximated to be isotropic. Provided the detector geometry remains fixed during a measurement, its angular dependence is inconsequential and acts only to change the numeric value of σ_{NRF} slightly from that tabulated for nuclear data. Equation S.8 can be generalized by rewriting the last two factors:

$$\exp[\dots] \left[\sigma_{\text{NRF}} \frac{\Omega}{4\pi} \right] \longrightarrow \frac{1}{4\pi} \int_{\Omega} \exp[\dots] \frac{d\sigma_{\text{NRF}}}{d\Omega} d\phi d\cos\theta$$

The treatment is also limited to cases where an excited state decay to the ground state, which are predominant for most NRF excitations. For instances where excitations decay through intermediate states, the value of σ_{NRF} needs to be calculated for the excited state i and final state j , with the substitution $\mu_{\text{NRF}} \rightarrow b\mu_{\text{NRF}}$ in the numerator of the linear term in Equation S.9, where $b = \sigma_{\text{NRF},ij}/\sigma_{\text{NRF},\text{total}}$ is the branching ratio for the particular transition. Furthermore, the value for $\mu_{nr}(E)$ going out of the foil will need to be evaluated at the energy of $E = E_i - E_j$. It should be noted that for the actinides of interest $E_i \sim \mathcal{O}(\text{MeV})$, $E_j \sim 40 \text{ keV}$ and therefore $E \approx E_i$.

4.2 Heterogenous device

The above calculations were done for a single isotope, and a single path through a homogeneous object and foil. For a wider, isotropic beam through a geometrically non-homogenous, single-isotope object $\rho(r + s\theta)$, and using the notational conventions from Section 2, we have:

$$\Phi_{\text{NRF}}(E) = \int_{\mathbb{S}^2} \left\{ A_{nr}(E, \theta) \exp(-\mu_{\text{NRF}} \int \rho(r + s\theta) ds) \times [1 - \exp(T\mu_{\text{NRF}})] \frac{\Omega(\theta)}{4\pi} \right\} d\theta \quad (\text{S.11})$$

where ρ is a volume mass density map that when line integrated over $r + s\theta$ yields an effective areal density analogous to D ; the previous definition of T is still valid if the foil is always normal to the ray in direction $\theta \in \mathbb{S}^2$; and the solid angle of the detector Ω is now a function of θ . For multi-isotope objects, the relevant parts of the above equations can be summed over all isotopes.

4.3 Non-invertibility of single-energy measurements

The bold variables in Eq. S.10 are unknown to the verifier but remain fixed between measurements of the candidate and template objects for a given projection. Specifically:

- \mathbf{X} is the areal density of the foil is fixed because the foil remains in place, but is unknown to the inspector.
- \mathbf{D} is the areal density along a projection of the object, is unknown, and needs to remain so in order to protect warhead design information.

Equation S.10 is a single equation with two unknowns. Any changes in \mathbf{D} will translate to a difference in the measured counts, enabling the detection of mismatched objects. However, absent knowledge of \mathbf{X} , it is impossible determine \mathbf{D} . Provided \mathbf{D} remains unchanged between measurements of a candidate and template warhead, which is approximately true if they match (see Section 7), then nothing about \mathbf{X} can be discerned between the two measurements for which \mathbf{X} is unchanged. Provided \mathbf{X} is then changed for each projection during which \mathbf{D} varies, all values of \mathbf{D} will be protected. The foil that provides \mathbf{X} therefore acts similarly to a one-time pad.

5 Monte Carlo Simulations

As a proof of concept, several Monte Carlo simulations of nuclear resonance fluorescence interactions in nominal object and hoax object geometries are carried out using the Geant4 toolkit [9]³. NRF interactions are modeled using an updated version of the G4NRF package [11] alongside the standard electromagnetic physics models. Each simulation was run for $\sim 10,000$ core-hours on Amazon's EC2 framework.

5.1 Bremsstrahlung interrogation beam

A bremsstrahlung beam was used to interrogate the object and reference foil. The spectrum is first generated in a standalone Geant4 simulation. A pencil beam of monoenergetic 2.7 MeV electrons impinges on a gold target of thickness $102 \mu\text{m}$ with a copper backing of thickness 20.1 mm. The subset of bremsstrahlung photons within 10° of the incident beam axis is then used to create the interrogation beam spectrum. This spectrum is smoothed using an interpolation routine to reduce the effects of discrete bin edges, then fed into the main Geant4 simulation as a pencil beam of photons directed towards the center of the warhead.

Because of computational limitations, the bremsstrahlung spectrum is not sampled directly. Although the integrated NRF cross sections (Eq. S.2) are relatively strong ($\int \sigma(E) dE \sim \mathcal{O}(10 \text{ eV}\cdot\text{b})$), the energy-dependent NRF cross sections are exceptionally narrow with Doppler-broadened widths of $\Delta \sim \mathcal{O}(1 \text{ eV})$. If sampling from a broad-spectrum bremsstrahlung spectrum, an NRF interaction would be a rare event because the probability that the random number generator (or the nature of the bremsstrahlung process) will choose a photon energy sufficiently close to the resonant energy to excite an NRF transition is small.

³Much of this section was written in parallel with a similar section in the SM thesis of one of the authors (JRV) [10].

Thus, the computational effort required for obtaining a statistically significant NRF signal using a naïve sampling of the full bremsstrahlung beam is thus enormous. This is a classic variance reduction problem that can be mitigated by importance sampling.

Instead of directly and inefficiently simulating the incident bremsstrahlung beam $\phi_0(E)$, an artificial sampling spectrum $s(E)$ that emphasizes NRF transitions of interest⁴ and de-emphasizes off-resonant photons is used. To adjust for this relative over- and under-sampling, respectively, each incident photon is tagged with a statistical weight factor $w(E) = p(E)/s(E)$, where $p(E)$ is the probability density function (PDF) of the true spectrum $\phi_0(E)$ and $s(E)$ is also normalized to unity. Plots of the $p(E)$ and $s(E)$ used in the simulations can be seen in Figure S.4. This weight is then propagated throughout the simulation to the secondary particles produced by each primary photon. To recover the output NRF spectrum that would correspond to sampling from $p(E)$ rather than $s(E)$, every event's contribution to the final recorded spectrum needs to be weighted by $w(E)$, where E is the initial energy of that event. This amounts to filling the final histogram not with one entry per signal photon, but with $w(E)$ entries per signal photon, where $w(E)$ is again the weight calculated from the *initial* photon energy. Since $s(E) > p(E)$ and thus $w(E) < 1$ for NRF photons of interest, this corrects for the oversampling of the NRF lines with respect to the true bremsstrahlung distribution.




Figure S.4: Smoothed Monte Carlo bremsstrahlung spectrum $p(E)$ (black) and sampling spectrum $s(E)$ (red) used in the simulations. The PDFs are normalized such that the sum rather than the true integral of their bin contents is each unity, but the distinction is immaterial when taking the ratio $w(E) = p(E)/s(E)$ as long as the distributions have identical bin widths.

⁴Though each of the four isotopes (^{238}U , ^{235}U , ^{240}Pu , and ^{239}Pu) has ten or so experimentally-observed lines, the ENSDF-derived database used by G4NRF lacks data for level spins J_r and widths Γ_r for many of the lines of the latter three isotopes. In some cases, reasonable estimates of these values can be made such that the observed integrated cross sections are replicated in the simulation. In other cases, such as the pair of ^{235}U transitions at 1769 and 1815 keV, this is not so, and the NRF signal of these incompletely-understood states is excluded from the model.

5.2 Simulation geometries

For the purposes of simulation, a nominal nuclear warhead geometry is constructed from an approximate characterization of a Soviet thermonuclear tactical warhead (hereafter the “Black Sea” warhead) as estimated from information in the public domain [12]. While the precise dimensions of the warhead are classified and thus unknown, the purpose of this work is to achieve a proof of concept, for which the approximate dimensions are sufficient. Using simplified isotopic models of the special nuclear material and the midpoints of the ranges of dimensions quoted in the paper, the Black Sea geometry simulated in Geant4 consists of three concentric shells of weapons-grade plutonium (WGPu), high explosive (HE), and highly-enriched uranium (HEU). The assumption of concentricity has been imposed.

Several hoax configurations are also considered. Two isotopic hoaxes are constructed by exchanging the WGPu for either fuel-grade plutonium⁵ (FGPu) or pure U-238. The latter being a close analogue of depleted uranium. These two hoaxes demonstrate the ability of the system to detect material nonconformity in a candidate object. A geometric hoax (see definition in Section 6.2.1) is constructed by arranging two sets of rectangular slabs to mimic the sequence of attenuations in the Black Sea warhead. The geometric hoax demonstrates the ability of the system to detect geometric nonconformity in the candidate object.

The reference foil is a homogenized sheet of U-235, U-238, Pu-239, and Pu-240 (equal parts by mass), 2 cm thick with an overall density of $\rho = 19 \text{ g/cm}^3$.

Because of computational constraints we do not simulate energy deposition in a realistic detector geometry, which would result in most of the computed photons being discarded. Instead, we tally all photons with energy $1.7 \text{ MeV} \leq E_\gamma \leq 2.7 \text{ MeV}$ exiting the foil with a backwards angle of $\theta > 3\pi/4$ (i.e. the backward opening cone with solid angle 1.8 steradians) relative to the forward beam direction. To account for detector effects, the resulting photon counts are then multiplied by a realistic high purity germanium (HPGe) detector’s geometric and intrinsic peak efficiencies.

Dimensions and material compositions of each geometry are summarized in Table S.1, and diagrams of the Black Sea warhead, geometric hoax, and reference foils are shown in Figure S.5. The full set of Monte Carlo results is shown in Figs. S.6–S.11 and described in detail in Section 5.3. The plot of the Black Sea NRF spectrum alone is shown in Figure 2 of the main text.

⁵Since 1976, the U.S. Department of Energy has defined Fuel Grade Plutonium as an intermediary between Weapons-Grade Plutonium and Reactor-Grade Plutonium. Fuel Grade Plutonium is defined as plutonium with between 7% and 19% Pu-240. Nuclear weapons made from Fuel Grade Plutonium have been successfully detonated. See: <http://permanent.access.gpo.gov/websites/osti.gov/www.osti.gov/html/osti/opennet/document/press/pc29.html>

Object	Material layers	Composition (by mass percent)	Dimensions [cm]
Black Sea template	WGPu	94% Pu-239, 6% Pu-240	$(r_{\text{in}}, r_{\text{out}}) = (6.27, 6.7)$
	HMX	3% H, 16% C, 38% N, 43% O	$(r_{\text{in}}, r_{\text{out}}) = (6.7, 13.2)$
	HEU	95% U-235, 5% U-238	$(r_{\text{in}}, r_{\text{out}}) = (13.2, 13.45)$
Black Sea, with WGPu \rightarrow U-238 substitution hoax	U-238	100% U-238	$(r_{\text{in}}, r_{\text{out}}) = (6.27, 6.7)$
	HMX	3% H, 16% C, 38% N, 43% O	$(r_{\text{in}}, r_{\text{out}}) = (6.7, 13.2)$
	HEU	95% U-235, 5% U-238	$(r_{\text{in}}, r_{\text{out}}) = (13.2, 13.45)$
Black Sea, with WGPu \rightarrow FGPu substitution hoax	FGPu	86% Pu-239, 14% Pu-240	$(r_{\text{in}}, r_{\text{out}}) = (6.27, 6.7)$
	HMX	3% H, 16% C, 38% N, 43% O	$(r_{\text{in}}, r_{\text{out}}) = (6.7, 13.2)$
	HEU	95% U-235, 5% U-238	$(r_{\text{in}}, r_{\text{out}}) = (13.2, 13.45)$
Geometric hoax (see Figure S.5)	WGPu	94% Pu-239, 6% Pu-240	$\ell = 2 \times 0.43$
	HMX	3% H, 16% C, 38% N, 43% O	$\ell = 2 \times 6.50$
	HEU	95% U-235, 5% U-238	$\ell = 2 \times 0.25$
Reference foil	Homogenous	25% each U-235, U-238, Pu-239, Pu-240	$\ell = 2$

Table S.1: Object, hoax, and reference foil geometries simulated in Geant4, using weapons-grade plutonium (WGPu), fuel-grade plutonium (FGPu), high explosive (HMX formula), highly-enriched uranium (HEU), U-238, and a four-isotope composite. The pair $(r_{\text{in}}, r_{\text{out}})$ denotes the inner and outer radii of the spherical shell, whereas ℓ denotes the sheet thickness along the beam axis. When elements are given rather than isotopes, the isotopic composition of each element is assumed to be natural.




Figure S.5: Graphical rendering of the Geant4 models of the Black Sea warhead (shown in cutaway, left) and the dual-slab geometric hoax (right). The plutonium layers are shown in orange, high explosive in green, uranium in blue, and reference foil in gray. The simulation only registers the events corresponding to photons emitted from the foil at $\theta \geq 135^\circ$ angles, relative to the beam axis.

5.3 Monte Carlo results

Figures S.6 to S.11 show the Monte Carlo simulated photon spectra in the angular acceptance region $\theta > 3\pi/4$ for various scenarios. In each plot, two weighted spectra resulting from $N_{\gamma}^{\text{sim}} = 8 \times 10^{11}$ incident photons are compared; the black curve is fixed to be the spectrum from the authentic Black Sea template, while the red curve is the spectrum of the particular candidate object. The error bars on each energy bin are given by the square root of the sum of the squares of weights $\text{err}_E = \sqrt{\sum_i w_i^2(E)}$ in that bin, where $w_i(E)$ is the weight of the i -th event contributing to bin E , as determined in section 5.1. For the oversampled NRF lines this results in much smaller errors than simply the square root of the counts, as would be the case for simple Poisson errors. These errors represent the uncertainty on the expected signal strength as determined by the

simulation, rather than the uncertainty due to counting statistics that would be seen in a real experiment. When calculating the discrepancies between authentic and candidate objects the value of err_E is replaced with the Poisson error, i.e. \sqrt{N} where $N = \sum_i w_i(E)$ is just the bin content, as outlined in Section 5.4.

In Figure S.6, two authentic template objects are compared. Computationally, this amounts to simulating the Black Sea warhead twice, with different random number generator seeds used to sample the interrogation spectrum. The two spectra agree well within error bars, as expected.

In Figure S.7, the first hoax scenario is considered: the hosts use the Black Sea warhead geometry but replace the WGPu component with pure U-238. This type of hoax will successfully spoof neutron radiography systems [7] but will still be detected in NRF measurements (see Table S.2). By Eq. S.9, the presence of additional U-238 in the candidate object strongly *decreases* the U-238 NRF signal seen by the detector, since the object becomes a stronger filter of the lines corresponding to U-238. The corresponding loss of Pu-239 and Pu-240 in the candidate object *increases* the signal from these two isotopes when compared to the authentic template. Note that the continuum underlying the NRF lines also increases as the $Z = 94$ of the WGPu is replaced by the $Z = 92$ of the U-238; the continuum is a result of the photoelectron bremsstrahlung process which scales strongly with Z and thus is attenuated less in the hoax object.

In Figure S.8, the WGPu is instead replaced with FGpu. In this substitution, the mass fraction of Pu-239 in the plutonium shell decreases from 94% to 86% (a change by a factor of 0.91) while the mass fraction of Pu-240 changes from 6% to 14% (a change by a factor of 2.3). This substantial increase in Pu-240 content manifests itself as a noticeable reduction in the Pu-240 lines of the candidate object.

In Figure S.9, the geometric hoax (see Figure S.5) that mimics the sequence of attenuations in the authentic Black Sea warhead is interrogated. By design, it is undetectable when perfectly aligned with the beam (rotation of 0°). The two spectra agree well within statistics, indicating a successful hoax at that particular angle.

In Figure S.10, the geometric hoax of Figure S.9 is rotated by 10° , such that more of the hoax material lies in the beam path. Most NRF lines considered show a weaker signal in the rotated hoax as compared to the rotated authentic object.

For completeness, in Figure S.11, the geometric hoax is rotated to 30° . The further increase in the amount of material in the beam path causes the final signal to be even more strongly attenuated than in Figure S.10. Now, both the entire set of NRF lines and the continuum are substantially reduced in the candidate spectrum. In Section 5.4, it will be shown however that a rotation by 10° is sufficient to detect this particular hoax with 4σ confidence.

Scenario	^{238}U	^{235}U	^{239}Pu	^{240}Pu	$1-\text{P}_{\text{alarm}}$ (4σ threshold)	Figure (for reference)
Template vs authentic	1.41	0.454	1.47	1.69	0.98	S.6
Template vs WGPu \rightarrow U-238	127	3.49	14.0	1.48	0.0	S.7
Template vs WGPu \rightarrow FGpu	3.75	1.93	1.49	7.16	4.6×10^{-4}	S.8
Template vs geometric hoax 0° rotation (perfect alignment)	1.81	0.96	0.747	0.320	0.98	S.9
Template vs geometric hoax, 10° rotation	8.83	4.40	0.0796	0.964	2.3×10^{-7}	S.10
Template vs geometric hoax, 30° rotation	76.7	25.3	14.2	45.6	0.00	S.11

Table S.2: Detection probabilities for the strongest lines of the four isotopes in various hoax scenarios, as measured in units of σ . The values are computed for the strongest simulated NRF line of each isotope, using a 21-second run as described in the text. The “ $1-\text{P}_{\text{alarm}}$ ” column indicates the probability that none of the four comparisons cross the 4σ threshold. The “Figure” column indicates the figure containing the Monte Carlo results that are scaled and then compared via Eq. S.14 to compute the values in the isotope columns. The specific lines used are 2176 keV (U-238); 1957 keV (U-235); 2431 keV (Pu-239); and 2433 keV (Pu-240). The non-unity value for the $1-\text{P}_{\text{alarm}}$ in the template vs authentic scenario is only due to the limited sampling in the Monte Carlo simulation.




Figure S.6: Monte Carlo NRF spectra of the Black Sea control geometry (black curve) vs the identical Black Sea geometry with different Monte Carlo random seeds (red curve). The two spectra agree well within statistics. Lines corresponding to the four main NRF isotopes are marked by colored arrows. The errors are described in the text.




Figure S.7: Monte Carlo NRF spectra of the Black Sea control (black) vs the Black Sea geometry with WGPu replaced by U-238 (red). The U-238 lines (light blue arrows) are strongly diminished in the hoax, while the Pu-239 lines (green arrows) and Pu-240 lines (orange arrows) are enhanced.

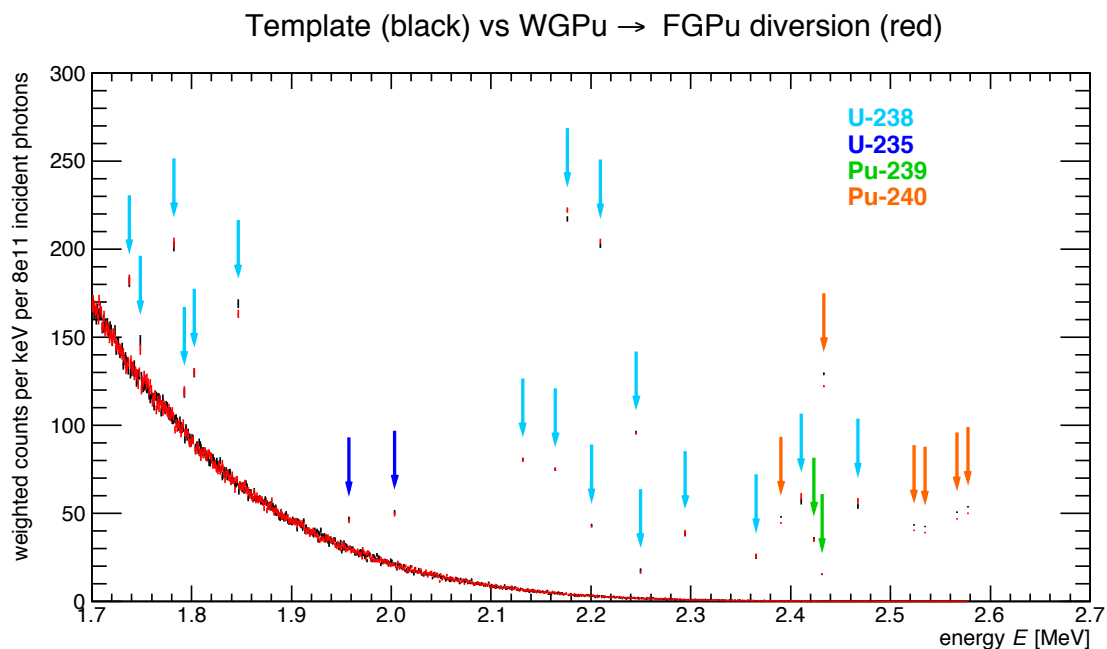


Figure S.8: Monte Carlo NRF spectra of the Black Sea control (black) vs the Black Sea geometry with WGPu replaced by FGPu (red). The Pu-240 lines (orange arrows) are diminished in the hoax, but the change in Pu-239 lines (green arrows) is not easily detectable with the statistics available.

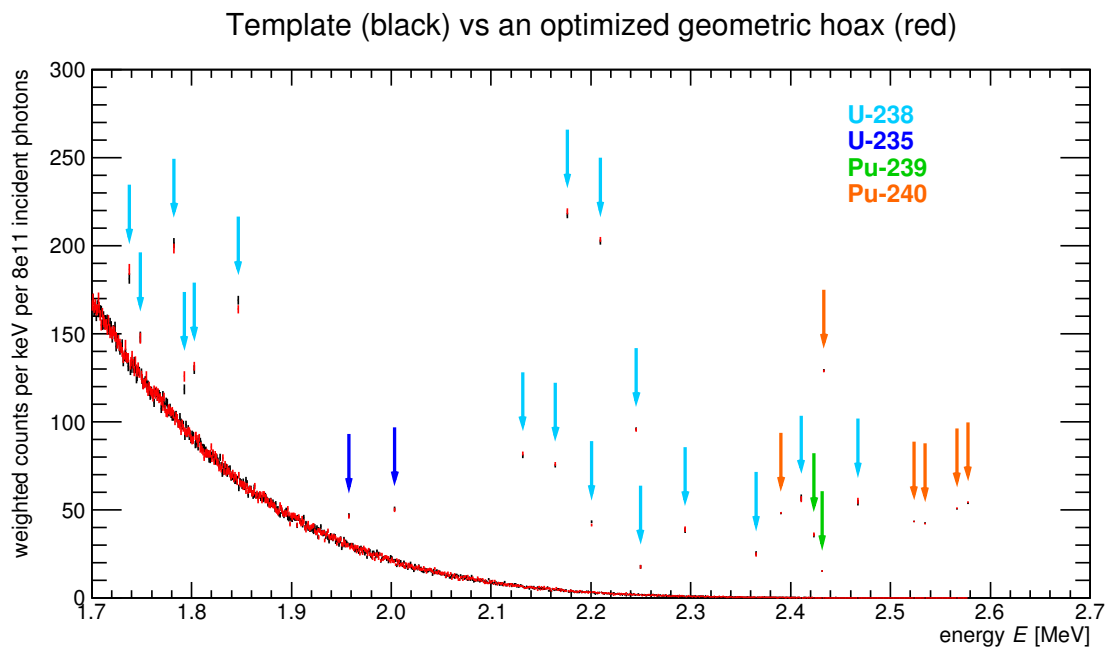


Figure S.9: Monte Carlo NRF spectra of the Black Sea control (black) vs the unrotated geometric dual-slab hoax (red). The two spectra agree well within statistics.




Figure S.10: Monte Carlo NRF spectra of the Black Sea control (black) vs the geometric dual-slab hoax rotated by 10 degrees (red). The two spectra no longer agree well within statistics.




Figure S.11: Monte Carlo NRF spectra of the Black Sea control (black) vs the geometric dual-slab hoax rotated by 30 degrees (red). The two spectra disagree strongly in both the NRF lines and the continuum.

5.4 Extrapolation to fixed measurement times

It remains now to extrapolate from the $N_\gamma^{\text{sim}} = 8 \times 10^{11}$ incident photons used for the weighted spectra in Section 5.3 to the number of photons N_γ^{exp} produced in a realistic experiment. The total number of electrons injected into an accelerator beam over a time t is $N_e = I_{\text{beam}}t/q$, where I_{beam} is the beam current (in Amperes) and $q = 1.602 \times 10^{-19}$ C is the (absolute) electron charge. The total number of photons across the entire bremsstrahlung energy spectrum is then just $N_\gamma^{\text{exp}} = f_{\text{conv}}N_e$, where f_{conv} is a dimensionless conversion factor, obtained from the Geant4 bremsstrahlung simulations, that gives the number of bremsstrahlung photons per electron incident on the radiator. For the angular acceptance of 10° used here, f_{conv} has been determined to be $f_{\text{conv}} \simeq 0.0144$, independent of the artificial sampling distribution $s(E)$. Altogether, the total number of incident photons N_γ^{exp} produced in a single accelerator run is

$$N_\gamma^{\text{exp}} = f_{\text{conv}}I_{\text{beam}}t/q, \quad (\text{S.12})$$

which amounts to a rate of $N_\gamma^{\text{exp}}/1\text{ s} = 3.6 \times 10^{14}$ incident photons per second with a proposed beam current of 4 mA [13]. To extrapolate from the Monte Carlo results of Figures S.6–S.11 to the counts and therefore the beam time needed for a realistic experimental setup, the spectra are scaled linearly by a factor

$$k \equiv \epsilon_{\text{int}}\epsilon_{\text{geom}}M_{\text{det}}N_\gamma^{\text{exp}}/N_\gamma^{\text{sim}}, \quad (\text{S.13})$$

where the geometric and intrinsic efficiencies of a number M_{det} of realistic HPGe detectors have also been taken into account. For a cylindrical detector of radius 5 cm positioned 1 m from the foil, the geometric efficiency is $\epsilon_{\text{geom}} \simeq 0.008/1.8/\text{detector}$ (to adjust from the 1.8 steradians of the $\theta > 3\pi/4$ signal region to the realistic solid angle) and the intrinsic efficiency is $\epsilon_{\text{int}} \simeq 0.20$ [14]. The value of this scale factor is set to be $k \simeq 254$, assuming $M_{\text{det}} = 30$ such detectors, as for these values, the significances in Table 2 will be reached with $t = 21$ s. Note that these results are obtained for a foil thickness of $X = 2$ cm; if a thinner foil is required for enhanced information security (see Section 7), fewer counts per second will be generated in the foil and the measurement times will increase roughly in accordance with Figure S.12. Furthermore, upon scaling, the Monte Carlo error bars are replaced by the Poisson error bars that would be seen in a real experiment. Explicitly, an NRF line with n counts and error bars of $\pm\delta n$ in the Monte Carlo simulation (where $\delta n \neq \sqrt{n}$ in general) is scaled to a value of kn with error bars of $\pm\sqrt{kn}$.




Figure S.12: Plot of the ratio of the time required to achieve a desired NRF signal strength in the $E_r = 2.176 \text{ MeV}$ uranium-238 line with an arbitrary total foil thickness X versus the time required with $X = 2 \text{ cm}$. For illustrative purposes, ratios are calculated via a modified version of Eq. S.9 at $\theta = 0$ using only uranium-238 data with the μ_{NRF} terms divided by 4 to more closely model the NRF response of the simulated composite foil. For the $E_r = 2.176 \text{ MeV}$ decay to the ground state, the NRF decay cross section is $\sigma_{\text{NRF}}(E_r) = 31.5 \text{ b}$ and the non-resonant cross section is $\sigma_{\text{nr}}(E_r) = 18.8 \text{ b}$. For $X \ll 1 \text{ cm}$, the thin foil limit of Eq. S.9 is reached, in which the NRF counts are directly proportional to X and therefore the time required to obtain a fixed number of counts is inversely proportional to X . For $X \gg 1 \text{ cm}$, the saturation limit of Eq. S.9 is reached, and the measurement time does not improve substantially.

6 Statistical Tests of Identity

6.1 Statistics of single projection comparisons

The NRF lines in spectra from a candidate object are compared to the NRF lines in the spectra of the template object to determine if the two objects match. The deviation in the strength of the NRF lines can be expressed in units of σ by

$$\Delta = \frac{c_{\text{can}} - c_{\text{tem}}}{\sqrt{\sigma_{\text{can}}^2 + \sigma_{\text{tem}}^2}}, \quad (\text{S.14})$$

where c_{can} and c_{tem} denote the counts (or scaled counts if using simulated data) in the NRF energy bin of interest for the candidate and authentic template object, respectively, and σ_{can} and σ_{tem} are the corresponding errors. For the lines with little or no continuum background, the errors are approximately the square root of counts. For the hoax scenarios considered in the main text, the differences between the hoax and template for at least one line always exceeded 7σ in less than 21 array-seconds, much greater than the 4.2σ detection threshold required to achieve a false alarm rate of 1 per 10,000 warheads.

6.2 Statistics of sampling multiple projections

The measurement of all NRF lines observed in a single exposure with a single foil is termed a *projection*. Each projection is indexed by a tuple (r, θ) , where r is the location of the x-ray source and θ is the direction

of the emitted beam. Each projection is a sample from the space of all possible spatial projections of the \mathcal{K} transform of the warhead's geometry, as described in Section 2. Because the transform exhibits global uniqueness (see proof in Section 2.2), sampling all possible projections will fully determine the underlying geometry of the object being measured and will provide a theoretically perfect measure of the warhead's authenticity. In practice, it is possible to determine if the object's geometry matches that of the template with a high confidence using only a few projections, provided each projection is randomly selected. This section outlines a method for estimating a lower bound of on the confidence for detecting a geometric hoax after N projections.

6.2.1 Definition of a geometric hoax

An object with the appropriate isotopic mix able to appear as a match for one or more projections, but which does not match the template's geometry, is called a "geometric hoax." As described in the main text, it is possible to show that a geometric hoax can be constructed for any finite number of projections, although they become ever more complex and thus impractical to manufacture as the number of projections increases. The shapes of undetectable hoaxes $(f + g)$ for a finite set of H projections $(r, \theta)_j$ where $j = \{1, \dots, H\}$ can be derived from the null-space of the \mathcal{K} transform given an authentic object f :

$$\mathcal{N}_{\mathcal{K}} = \{g \mid \mathcal{K}(f + g)(r_j) = \mathcal{K}(f)(r_j) \forall (r, \theta)_j\} \quad (\text{S.15})$$

where g is the geometry of a "ghost" map that if added to f is undetectable. This null space has not been formally characterized and remains one of the open areas for research. However, our studies of \mathcal{K} transform suggest that a property of the Radon transform, on which the \mathcal{K} transform is based, applies to the \mathcal{K} transform as well; namely, that the number of vertices in a valid ghost g scales with H [15]. This property suggests that geometric hoax objects will rapidly become highly oscillatory and complex as H is increased. This is an essential finding as it significantly limits the search space to geometric hoaxes of low order which are equal or less difficult to manufacture than a real warhead; cheating by high-order geometric hoax would be more difficult than simply re-manufacturing a new true warhead. Note also that a geometric hoax does not permit the cheater to use nonconforming materials, only to arrange them in nonconforming geometries.

6.2.2 A discrete approximation for estimating the detectability of a geometric hoax

The probability of detecting the hoax at a hoax-functional projection $(r, \theta)_j$ is identically zero, and the probability of alarm at such a projection is α , where α is the false-alarm rate for the system. As the projection parameters deviate from any given $(r, \theta)_j$, the probability of alarming on the hoax increases as a smooth, continuous, but unknown function of $(f + g)$. This gives rise to a problem in geometric probability that is not yet tractable because the null space $\mathcal{N}_{\mathcal{K}}$ is not yet characterized. However, a lower bound on the detection rate can be achieved by rewriting the problem as a simpler one with discretized probabilities.

We divide the space into two domains: one domain consisting of all the neighborhoods around each of the H hoax-functional projections $(r, \theta)_j$, in other words $\{(r \pm \delta r, \theta \pm \delta \theta)_j \mid \forall j\}$. All projections falling in this domain are assumed to have alarm rate of the minimum in those neighborhoods, namely that of $(r, \theta)_j$, which is by definition α . Every projection external to this domain is assumed to have probability of alarm $(1 - \beta_j)$, which is the alarm rate defining the edge of each neighborhood j , and therefore also a lower bound on the probability of alarm for the entire space external to the neighborhoods; and where β_j is the false-pass rate (FPR) for a given hoax at a projection $(r \pm \delta r, \theta \pm \delta \theta)_j$.

The inferred system-wide false-pass rate for hoax objects B , after k randomly selected projections, is then distributed according to a multivariate hypergeometric distribution in k . The multivariate hypergeometric distribution is the sample-with-no-replacement version of the multinomial distribution, where sample-with-no-replacement means that previously chosen projections will not be repeated. The calculation of B can be illustrated as a decision tree, as shown in Figure S.13A. While it is simple to calculate the probability of any given path through this tree, the cumulative distribution function (CDF) for the entire tree cannot be expressed in simple, practical terms. Fortunately, the multinomial distribution, which arises when replacement is allowed, provides a close and bounding approximation to the hypergeometric distribution. The multinomial distribution is shown in decision tree in Figure S.13B. The relationship between these two distributions will be demonstrated here.




Figure S.13: Decision tree describing the protocol for $k = 1-3$ projections. There are N possible projections for the system, for which a geometric hoax will appear authentic for H projections. Some subset k projections are chosen using uniform random sampling, either [A] with no replacement; or [B] with replacement. A candidate object will only pass the test if all k projections pass.

The question of interest is how often a hoax warhead will be accepted by the proposed testing protocol: $B_{k,\text{nr}} = \text{Prob}[\text{hoax accepted, no replacement}]$. According to the formalism adopted in Section 6.2.2, a hoax object can be imaged from N projection positions, H of which are perfect replications of the equivalent projection of the authentic object. These are statistical comparisons and as such random fluctuations may affect the result. The per-projection false pass rate β_j is the proportion of unmatched projections which pass the test due to statistical fluctuations given the allowed false alarm rate α , where α is the proportion of matching projections which fail the test because of statistical fluctuations. The true pass rate is the fraction of matched projections which correctly pass and is $(1 - \alpha)$, whereas the true alarm rate is the fraction of hoaxes that fail the test for projection j and is $(1 - \beta_j)$.

The four branches moving to the right from each node in Figure S.13 correspond to choosing a matched or unmatched projection and then correctly or incorrectly assessing whether they match the authentic projection or not. B is calculated only for outcomes where all of the tests are passed, correctly or not, so the middle two branches of each node can be discarded. B is then given by the sum of the remaining terms in the k^{th} column of the decision tree. The k projections are chosen using uniform sampling, either with or without replacement, then B is described by Equations S.16 and S.17, respectively:

$$P_{k,r} = \sum_{j=0}^k \left[\frac{k! \left(\frac{H}{N}(1 - \alpha) \right)^j \left((1 - \frac{H}{N})\beta_j \right)^{k-j}}{j! (k-j)!} \right] = \prod_{j=1}^k \left[\frac{H}{N}(1 - \alpha) + (1 - \frac{H}{N})\beta_j \right], \quad (\text{S.16})$$

$$P_{k,\text{nr}} = \sum_{j=0}^k \left[\frac{\binom{H}{j} \binom{N-H}{k-j} (1 - \alpha)^j \beta_j^{k-j}}{\binom{N}{k}} \right]. \quad (\text{S.17})$$

Proving that $P_{k,\text{nr}} \leq P_{k,r}$ always has not been achieved for all k . However this statement can be trivially proven for small $k < 5$ by expansion of the series and also demonstrated numerically for all k and all values of H and N , assuming $N \geq H$. Given that a practical implementation of the verification system will have $k < 5$ and that the multivariate hypergeometric distribution has been previously proven to approach the multinomial distribution as $N \rightarrow \infty$, it was considered allowable to assume $P_{k,\text{nr}} \leq P_{k,r}$.

Figure S.14 shows numerical simulations of the system. The top two plots show that $P_{k,\text{nr}} \leq P_{k,r}$ is valid for a range of k and that the two are equal for small k . Importantly, as two move away from equality, the value of both probabilities decrease so the absolute error also decreases. The bottom two plots demonstrate how the system behaves for a fixed false alarm rate α and variable false-pass rate β_j and number of matching projections, as well as the percent relative error of the same. These plots highlight the difficulty of hoaxing the system, as the probability of success is small until $\beta_j \geq 0.5$ or $H/N \geq 0.6$, reflecting how closely a

successful hoax will need to match the authentic object. These probabilities are for only one warhead so a poor set of hoaxes will not work at all over hundreds of warheads and tests.

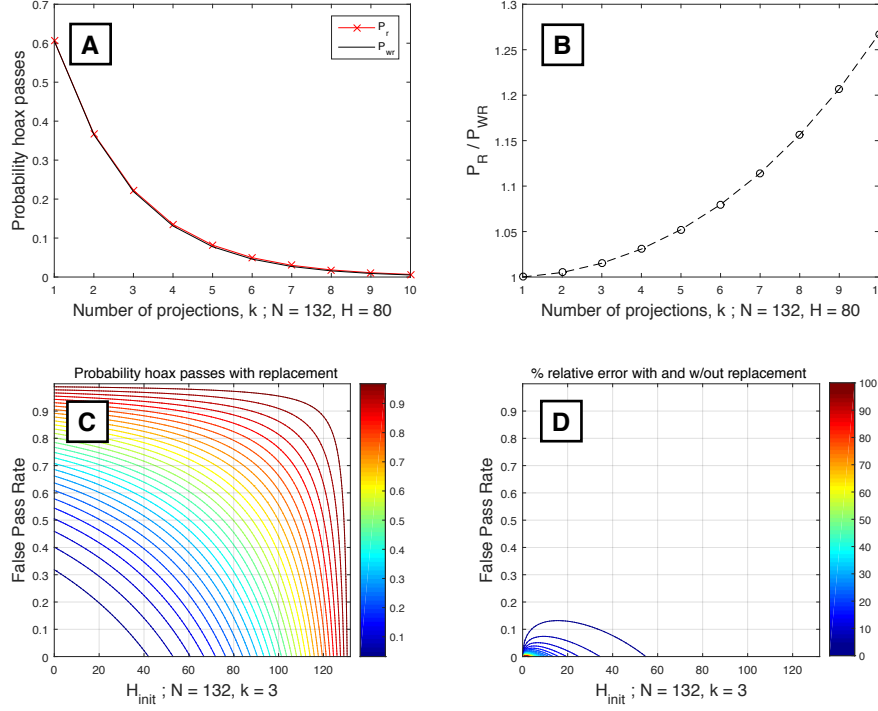


Figure S.14: Clockwise from top left: [A] Comparison of the probability that a hoax passes a test involving k projections, calculated with replacement ($B_{k,r}$) or with no replacement ($B_{k,nr}$). $N = 132$, $H = 80$, $\alpha \equiv 1 \times 10^{-4}$ and $\beta_j = 1 \times 10^{-4}$; [B] Ratio of $\frac{B_{k,r}}{B_{k,nr}}$ given k projections; [C] $B_{k,r}$ for various H and FPR, given $N = 132$ and $\alpha \equiv 1 \times 10^{-4}$; [D] The relative error between $B_{k,r}$ and $B_{k,nr}$ for $N = 132$ and $\alpha \equiv 1 \times 10^{-4}$.

7 Information Security

It is useful to recall here Equation S.11, which gives the counts observed at energy E from NRF:

$$\Phi_{\text{NRF}}(E) = \int_{\mathbb{S}^2} \left\{ A_{nr}(E, \theta) \exp(-\mu_{\text{NRF}} \int \rho(r + s\theta) ds) \times [1 - \exp(T\mu_{\text{NRF}})] \frac{\Omega(\theta)}{4\pi} \right\} d\theta \quad (\text{S.11})$$

Provided that each measurement at some given energy E is independent of another, then reconstruction of the test object $f = \rho(r, \theta)$ is strictly prevented by the following four properties:

1. The underdetermined nature of the system with respect to the measured counts in a single NRF bin alone. Given a measurement described by Φ , it is not possible to infer information about the object's composition (described by ρ) as the foil composition T is unknown (see discussion in Section 4.3).
2. The nonlinearity of the \mathcal{K} transform. If the inspector had knowledge of the foil (not normally allowed), the inspector could learn from a single projection the value of the “transform mass” $\tilde{m} \propto \iint \exp[-\int \rho(r, s\theta) ds] d^2\theta$, rather than the true mass $m = \iiint \rho_i(r, s\theta) ds d^2\theta$. There is no one-to-one correspondence between \tilde{m} and m , therefore the mass of any isotope present in the test object cannot be estimated from a single projection, even with knowledge of the foil.
3. The absence of spatial information in the zero-dimensional (single-pixel) \mathcal{K} transform means information about the geometry of the test object cannot be learned from a single projection, as one might

learn by evaluating shadows or edges in a spatially-resolving measurement (Section 2). Furthermore, if every measurement of the integrated density is adequately protected by a properly operating independent foil (property 1), then the shape can never be reconstructed regardless of the number of projections taken.

4. Finally, if very few projections (r, θ) are sampled from the transform space, the geometry of f cannot be reliably determined by inversion of $\mathcal{K}f$ even if the foil is known. In other words, there is inadequate sampling of the geometric information. The number of samples needed to determine the geometry of the object with confidence is ultimately one of geometric probability and depends on the shape of the warhead. However, the results presented here suggest that even one sample (r, θ) may be enough to determine the authenticity of the warhead with high confidence, which would never be sufficient to reconstruct the geometry.

In an ideal system, these information protection mechanisms work as layers of protection. Individually, however, one or more may break down given particularities of real-world implementation. Furthermore, certain plausible guesses could be made by the inspector, helping to fill in missing information. A strict proof of information security therefore requires the specific operating parameters of the system to be examined and may require that additional constraints be imposed on the measurement protocol to mitigate vulnerabilities. Several real-world limitations are discussed below, and solutions offered, to the extent they are understood. Ultimately, this remains an area of continuing research.

7.1 Information security in real systems

7.1.1 X-ray fluorescence

The process of x-ray (atomic) fluorescence may be stimulated in the foil by either highly down scattered photons from a monoenergetic interrogating beam, or if a bremsstrahlung beam is used. This fluorescence is generated by the excitation and decay of bound electrons around the nucleus. The emission spectrum contains characteristic peaks associated with the electron shell structure of the atom. These emissions, if observed by the detector, could reveal the relative elemental composition of the foil.

Solution. A 1 cm-thick tungsten plate can be inserted by the hosts along the line joining the foil and the detector, as illustrated in the main text. The tungsten will filter out essentially all photons in the < 200 keV range, thus removing x-ray fluorescence (XRF) peaks from the signal, which are in the ~ 10 keV range.

7.1.2 Nonmonochromatic NRF measurements

If a bremsstrahlung interrogating beam is used, or certain NRF transitions are excited, the single-energy assumption in property 1 is not upheld. Specifically,

- (a) some NRF transitions excited by a single incoming photon may release multiple, lower-energy photons that sum to the incoming photon's energy. If outgoing photons have significantly different energies, then significantly unequal attenuation of the outgoing photons by non-resonant process, such as Compton scattering, may be observed in the measured signal. This may be used to estimate the foil parameter t_i in Equation S.11, although with considerable uncertainty. Our initial studies suggest uncertainties in the range of 100% or greater may be expected over a wide range of foil thicknesses.
- (b) if the interrogating beam is produced from bremsstrahlung, as opposed to a monoenergetic process,⁶ multiple NRF transitions of a single isotope will be simultaneously excited. If the ratios of incident bremsstrahlung flux for two or more energy lines of a single isotope can be approximated (for example, if the resonant energies are clustered in an energy window small enough so that the bremsstrahlung spectrum is approximately flat, and the strength of the NRF cross sections are significantly different,

⁶Several tunable quasi-monochromatic sources are under development. Especially notable are those that use Compton backscattering of laser light from \sim GeV electron beams. Such sources as currently available in the form of major user facilities may not be practical as they do not yet have the intensities necessary for producing sufficient detector counts in a transmission NRF system [16]. Nevertheless, significant research efforts are underway to increase the photon fluxes and repetition rates at laser Compton backscattering sources.

then the ratio of outgoing fluxes can be used to infer the approximate thickness of the foil. Again, this reconstruction bears considerable uncertainty and is subject to a systematic error arising from the approximate assumption about the incoming bremsstrahlung flux.

Under the above two conditions, which hold only if certain controllable experimental conditions are chosen, the thickness of the foil may be estimated. This knowledge can then be used to reduce the number of unknowns in the system (property 1) and thereby estimate the transform mass \tilde{m} . However, property 2 will still prevent the inspector from learning the true mass m without knowledge of the geometry.

Potential Solutions. There are several ways to mitigate the vulnerabilities that arise from implementation choices that would result in the observation of multiple NRF lines:

- Measurements can be made one isotope at a time. The end-point energy of the bremsstrahlung beam can then be adjusted between measurements so that photon energies will only be sufficient to observe one NRF transition in the measured data. This solution precludes vulnerability (b) entirely, and eliminates vulnerability (a) provided the lowest statistically observable line is non-branching.
- A spatially non-homogenous foil of multiple regions, each region having a different isotopic composition, may be provided by the host but randomly oriented by the inspector. Each region of the foil will sample a different area of the warhead's projection and will weight that region by a different but unknown factor. The regional sampling is then averaged together by the single-pixel detector. The additional regions increase the number of unknowns in the system and/or add additional systematic errors to the reconstruction model. This solution should be useful for addressing both vulnerabilities (a) and (b).
- Additional material of each isotope in the form of plates or pucks may be added to the warhead container by the host. The amount and location of this material can be unknown to the inspector. Provided the geometry of the added material is identical in both the template and candidate objects, the system will be able to validate the weapon as expected, including the detection of hoax objects. However, any reconstruction of the transform mass \tilde{m} by the inspector will be contain an added effect associated with the unknown plate or puck. This effect can easily be of the same magnitude as the warhead, creating a significant error in the estimation. This solution should be useful for addressing both vulnerabilities (a) and (b).
- In a strategy similar to the multi-region foil, the spatial intensity and energy spectrum of an interrogating bremsstrahlung beam may be modulated by a spatially heterogeneous filter inserted prior to the test object. As with the multi-region foil, the coded aperture would be provided by the host and its structure would be unknown to the inspector, while its position in the beam would be selected by the inspector to ensure the host has not used it to facilitate cheating. This solution should be useful for addressing both vulnerabilities (a) and (b).
- If the thickness of the tungsten plate used to mitigate the x-ray fluorescence vulnerability (Section 7.1.1) is unknown to the inspector, it will add additional Compton scattering, disturbing the ratio of the lines observed and used in the reconstruction of vulnerability (a).
- If the bremsstrahlung can be shaped over a narrow energy spectrum, it will make the systematic errors of the reconstruction larger for vulnerability (b).
- If foils are thin enough, the error in reconstructing the integrated density will be large. Preliminary explorations suggest foils thinner than approximately 1 mm will produce large (\mathcal{O} 100%) errors.

7.1.3 Prior assumption of geometries in estimating mass

In some cases, the inspector might correctly guess the general geometry-class for an isotope in the warhead, in which case the nonlinearity of the \mathcal{K} listed as property 2 can be overcome. For example, it might be reasonable for the inspector to assume the plutonium was configured as a single spherical shell.

If the geometry has only a single parameter (e.g., a solid sphere with one radius, or a solid cube having one side length), and this geometry is correctly guessed by the inspector, and there is no inhomogeneity

in the sampling of the object induced by non-resonant attenuation, or filters, or multi-segment foils from the solutions proposed for addressing multi-energy vulnerabilities of Section 7.1.2, then an estimate of the transform mass \tilde{m} using the vulnerabilities identified in Section 7.1.2 can be used to estimate the true mass, m . Evidently, the host is unlikely to allow this condition to occur. Nevertheless, even under perfect conditions, a single projection does not confirm the underlying assumptions about the geometry required to make the estimate. Moreover, a higher order geometry, such as nested shells having two radii parameters, will require correspondingly more projections and corresponding reconstructions of \tilde{m} to estimate m .

Solution. This vulnerability can be mitigated or eliminated by invalidating the obvious assumptions. For example, the proposal that the host may insert into the container an additional plate or puck of material for those isotopes that are in easily-guessed geometries would render incorrect normal assumptions about the geometric distribution of that isotope.

8 Operational Considerations

The system proposed assumes an idealized operating environment. A practical system will require a considerable engineering effort and attention to protocol details to cope with the particularities of real weapons, real verification environments, mutual trust in system integrity, potentially unidentified national-security restrictions, and safety considerations. For these reasons, the development of a practical system will necessarily require the expertise of nuclear-weapon laboratories from participating governments. While these issues are beyond the scope of this paper, we have provisionally identified several considerations.

8.1 The measured object

We have described verification as occurring on abstract warheads. In reality, warheads may arrive at a dismantlement facility with various levels of packaging. The most intact configurations might include missile casings, bomb shells, or re-entry vehicles. Below that, configurations may include the electronic firing set for the nuclear weapon, deuterium-tritium reservoirs, neutron generators, weak- and strong-link safety systems, and the physics package. While the proposed method should work in principle at any level of disassembly, the more fully assembled a weapon, the longer the measurements will need to be to compensate for additional attenuation of the interrogating beam. A potential further reason for partial disassembly may be to eliminate geometric or chemical heterogeneity between weapons caused by retrofits or service campaigns.

To cope with these issues, it may be possible to disassemble weapons to an agreed level in a sealed facility. Afterward, the inspector can verify that the facility contains only treaty-accountable items or insensitive packaging materials of an inspectable nature. This process would reduce the significance of the verification finding: instead of determining that a full nuclear weapon was being dismantled, all that is verified is that a subassembly of a nuclear weapon is being dismantled. Provided the subassembly retains enough strategic significance, however, this should be acceptable. That said, a state might have surplus of weapon components at various levels of assembly, possibly including nearly complete “dud” weapons that failed quality inspection. The significance of those components will have to be understood if verification of partially disassembled weapons is permitted or required.

8.2 Dose

The measurement involves irradiating the warhead with ionizing radiation, which will give the warhead a dose that may result in microstructural material damage. If a bremsstrahlung source is used, that damage may be up to a million times larger than would occur if a monoenergetic interrogating beam were used. We do not know if dose is a relevant consideration. Provided there are no safety issues arising from dose, damage to sensitive components such as electronics should not be an important consideration provided the template and test warhead are both being dismantled. Nevertheless, we provide here an estimate of the dose for the reader’s consideration.

To estimate the dose to a warhead from a bremsstrahlung beam, we modified the Geant4 simulation described in Section 5. The tracking cuts that killed low-energy photons for computational speed were disabled and the sampling distribution $s(E)$ was extended to cover the entire 0–2.7 MeV energy range, as

low-energy photons contribute significantly and nonlinearly to the total dose. Using the `sensitiveDetector` functionality of Geant4, the total weighted dose to the entire outer uranium layer (of mass 10.4 kg) was computed as 3.2×10^{-9} Gy per 1×10^6 incident photons.

To estimate an upper bound on the *local* dose at the point where a narrow bremsstrahlung beam impinges on the warhead, we assume the shell of uranium is placed 10 cm from the bremsstrahlung radiator, which is the vertex of a conical beam with a 10° half-angle, giving a beam-spot radius of 1.76 cm at the surface of the uranium. Since the illuminated spot size is therefore small compared to the 13.45 cm outer radius of the uranium shell (see Table S.1), the segment of irradiated uranium can be approximated as a disk of radius 1.76 cm, and a height equal to the shell thickness (0.25 cm), which has a mass of 46.4 g. The initial dose value of 3.2×10^{-9} Gy is therefore scaled by a factor of $(10.4 \times 10^3 \text{ g})/(46.4 \text{ g}) \simeq 224$. Finally, to extrapolate from the 1×10^6 incident photons to the full measurement of $254 \times 8 \times 10^{11}$ photons, the dose is scaled as in Section 5.4 by a factor of $254 \times (8 \times 10^{11}/1 \times 10^6)$ to give an approximate value for the local dose of ~ 145 Gy for the 84 mC integrated beam charge used in the 21 s measurement. The dose is reduced by increasing the distance between the vertex of bremsstrahlung emission cone and the weapon surface: for example, by doubling the distance to 20 cm the dose will be further reduced to 36 Gy. Such doses, while large, are within the tolerances of modern radiation-hardened electronic devices [17, 18].

The heating of the warhead due to dose can be estimated by considering the power of the interrogating beam. The simulations in this work were based on a 2.7 MeV electron beam of 4 mA current. The kinetic power of the electron beam is then 11 kW. Simulations of the bremsstrahlung process show that only 0.2% of the energy of the electrons is radiated as photons. Thus the power of the bremsstrahlung beam is no more than 22 W. Since only part of the bremsstrahlung beam is absorbed by the weapon, the heating rate for the weapon is necessarily less than 22 W, which is distributed over more than 10 kg of mass. Given a combined uranium and plutonium mass of more than 10 kg, and noting the specific heat of these metals is approximately 120 kJ/kg/K, the 21 s of heating at 22 W would result in an average temperature increase of less than 1 K.

8.3 Electronics and count rates

Practical, non-idealized detector systems are limited in their maximum allowable count rates due to their intrinsic charge or light collection times or due the pulse shaping times required for electronic signal processing. In particular, commercial HPGe detectors can reliably operate at average count rates of up to $\sim 10^5$ cps with only modest degradation of resolution. To estimate the count rates in the proposed system, another simulation similar to that of Section 8.2 was carried out. An incident beam of 8×10^7 photons resulted in a total of 4.7×10^3 counts over the entire energy spectrum, corresponding to 1.2×10^{10} counts after scaling to the full measurement of $254 \times 8 \times 10^{11}$ photons. Split over 30 detectors operating for 21 s, this results in an average count rate of $\sim 20 \times 10^6$ cps in each detector, a factor of 200 larger than the practical upper limit. These total count rates can be reduced substantially by several means:

- As described in Section 7.1.2, thinner foils may be used to enhance the information security of the system. This has the dual effect of reducing the count rate by a factor of $\mathcal{O}(10)$, depending on the thickness of the foil—see Figure S.12.
- As described in Section 7.1.1, a tungsten filter may be placed between the foil and the detector to preferentially attenuate low-energy characteristic x-ray fluorescence photons that could reveal the elemental composition of the foil. Since low-energy photons comprise the bulk of the 20×10^6 cps count rate, the tungsten filter has the dual effect of substantially reducing the total count rate. In our simulations, a 1 cm tungsten filter reduces the total count rate by a factor of 10 while only reducing the count rate in the 2.176 MeV uranium-238 line by a factor of 2; for a 2.15 cm filter, the reduction factors are $100\times$ and $6\times$, respectively.
- Segmentation of the active detector volume, whether by electronic retrofits or by using a larger number of smaller-diameter detectors, will reduce the count rate per electronics channel without a corresponding increase in measurement time.
- Finally, it is always possible to reduce beam currents and increase measurement time.

8.4 System integrity

The system described in this work is assumed to operate according to the protocol and concept of operations outlined in the paper. Extra procedures may need to be established to assure the inspector that the critical parts of the system operate as designed and are not being manipulated to pass hoax objects. While the system is comprised of straightforward physical objects, some amount of advanced inspection or supervision may be needed, and this will bring with additional operational and procedural demands given the national-security context. These considerations need to be considered at the time the system is designed because it may be possible that a particular technology or operation may not be procedurally acceptable. For example, while a monochromatic photon source would provide stronger information security than the simpler bremsstrahlung source, the operation of monochromatic systems will be much harder to authenticate against tampering by the host.

References

- [1] U. Kneissl, H.H. Pitz, and A. Zilges. Investigation of nuclear structure by resonance fluorescence scattering. *Prog. Part. Nucl. Phys.*, 37:349–433, 1996.
- [2] R.F. Casten. Nuclear structure from a simple perspective. *Oxford University Press, Oxford*, 1990.
- [3] F.R. Metzger. Resonance fluorescence in nuclei. *Prog. Nucl. Phys.*, 7, 1959.
- [4] B.J. Quiter, B.A. Ludewigt, V.V. Mozin, C. Wilson, and S. Korbly. Transmission nuclear resonance fluorescence measurements of ^{238}U in thick targets. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 269(10):1130 – 1139, 2011.
- [5] W. Bertozzi, J.A. Caggiano, W.K. Hensley, M.S. Johnson, S.E. Korbly, R.J. Ledoux, D.P. McNabb, E.B. Norman, W.H. Park, and G.A. Warren. Nuclear resonance fluorescence excitations near 2 MeV in ^{235}U and ^{239}Pu . *Phys. Rev. C*, 78:041601, Oct 2008.
- [6] B. J. Quiter, T. Laplace, B. A. Ludewigt, S. D. Ambers, B. L. Goldblum, S. Korbly, C. Hicks, and C. Wilson. Nuclear resonance fluorescence in ^{240}Pu . *Phys. Rev. C*, 86:034307, Sep 2012.
- [7] A. Glaser, B. Barak, and R.J. Goldston. A zero-knowledge protocol for nuclear warhead verification. *Nature*, 510(7506):497–502, 06 2014.
- [8] S. Helgason. *Integral Geometry and Radon Transforms*. Springer, 2010.
- [9] S. Agostinelli et al. GEANT4: A simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.
- [10] Jayson R Vavrek. Monte Carlo simulations of a physical cryptographic warhead verification protocol using nuclear resonance fluorescence. Master’s thesis, Massachusetts Institute of Technology, 2016.
- [11] David V Jordan and Glen A Warren. Simulation of nuclear resonance fluorescence in Geant4. In *Nuclear Science Symposium Conference Record, 2007. NSS ’07. IEEE*, 2007.
- [12] Steve Fetter, Thomas B Cochran, Lee Grodzins, Harvey L Lynch, and Martin S Zucker. Gamma-ray measurements of a Soviet cruise-missile warhead. *Science*, 248(4957):828–834, 1990.
- [13] IBA Industrial. Rhodotron[®] E-beam accelerator: Capacity on demand. Technical report.
- [14] Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, Inc., 4 edition, 2010.
- [15] A.K. Louis. Ghosts in tomography: The null space of the radon transform. *Mathematical Methods in the Applied Sciences*, 3(1):1–10, 1981.
- [16] Henry R. Weller, Mohammad W. Ahmed, Haiyan Gao, Werner Tornow, Ying K. Wu, Moshe Gai, and Rory Miskimen. Research opportunities at the upgraded HI γ S facility. *Prog. Part. Nucl. Phys.*, 62:257–303, 2009.
- [17] Kari Leppl, Kirgiz Ssr, and Raimo Verkasalo. Protection of instrument control computers against soft and hard errors and cosmic ray effects. In *International Seminar on Space Scientific Engineering*, 1989.
- [18] David R Meyers. A brief discussion of radiation hardening of CMOS microelectronics, 1998. Retrieved from <http://www.osti.gov/scitech/servlets/purl/2453> on May 8.